

A Semantic Analysis of Wireless Network Security Protocols


Massimo Merro

(joint work with Damiano Macedonio)

Dipartimento di Informatica
Università degli Studi di Verona - Italy

NFM 2012 - Norfolk, April 3-5 2012

Wireless Networks (Ad hoc, Sensor, Vehicular, Mesh networks ...)



- **Some challenging features:**

- No fixed infrastructure
- Radio frequency channels
- Half-duplex channels
- Local broadcast
- Multi-hop communication
- **High vulnerability**

Attacking and securing wireless networks

- In a wireless network an attacker may:
 - compromise a node (node subversion)
 - alter data integrity
 - eavesdrop on messages
 - inject fake messages
 - waste network resources
 - etc
- Designing security protocols for wireless networks requires a deep understanding of their resource limitations (Processor, Memory, Battery power, etc)

A process algebraic approach to model wireless networks

Assumptions:

- **Synchronisation:** all nodes are synchronised using a *clock-correction* synchronisation protocol (this implies network connectivity)
- **Time:** proceeds in *discrete steps*; a global clock is supposed to be updated whenever all nodes agree on this, by synchronising on a special action σ ([Hennessy and Regan 1995])
- **Fictitious clock approach:** data transmission is assumed to take no time. This is reasonable if the actual time of transmission is negligible with respect to our time intervals
- **Nondeterminism:** untimed activities among nodes occur nondeterministically
- **Mobility:** Our nodes are stationary (as in most sensor networks); communication and node mobility are orthogonal concepts

The Syntax

Networks :

M, N	$::=$	$\mathbf{0}$	empty network
	$ $	$M_1 \mid M_2$	parallel composition
	$ $	$n[P]^\nu$	node (ν = set of neighbours of n)

Processes :

P, Q, R	$::=$	nil	termination
	$ $	$\sigma.P$	sleep
	$ $	$!\langle u \rangle.P$	broadcast
	$ $	$[?(x).P]Q$	receiver with timeout
	$ $	$[\sum_{i \in I} \tau.P_i]Q$	internal choice with timeout
	$ $	$[u_1 \dots u_n \vdash_r x]P; Q$	deduction

- The calculus is parametric wrt to a given *decidable* inference system

Labelled Transition Semantics (some rules)

$$(\text{Snd}) \frac{-}{m[!\langle v \rangle.P]^\nu \xrightarrow{m!v \triangleright \nu} m[P]^\nu}$$

$$(\text{Rcv}) \frac{m \in \nu}{n[[?(x).P]Q]^\nu \xrightarrow{m?v} n[\{^v/x\}P]^\nu}$$

$$(\text{Bcast}) \frac{M \xrightarrow{m!v \triangleright \nu} M' \quad N \xrightarrow{m?v} N' \quad \mu := \nu \setminus \text{nds}(N)}{M \mid N \xrightarrow{m!v \triangleright \mu} M' \mid N'}$$

$$(\text{Sleep}) \frac{-}{n[\sigma.P]^\nu \xrightarrow{\sigma} n[P]^\nu}$$

$$(\text{Timeout}) \frac{-}{n[[?(x).P]Q]^\nu \xrightarrow{\sigma} n[Q]^\nu}$$

$$(\text{TimeSync}) \frac{M \xrightarrow{\sigma} M' \quad N \xrightarrow{\sigma} N'}{M \mid N \xrightarrow{\sigma} M' \mid N'}$$

Simulation theory

We are interested in a **weak semantics** which abstracts over internal actions, $\xrightarrow{\tau}$

Weak transitions

They are defined as usual:

- $\hat{\alpha} \stackrel{\text{def}}{=} \xrightarrow{\tau^*} \xrightarrow{\alpha} \xrightarrow{\tau^*}$, if $\alpha \neq \tau$
- $\hat{\tau} \stackrel{\text{def}}{=} \xrightarrow{\tau^*}$

Definition: Similarity

- $M \lesssim N$ if $M \xrightarrow{\alpha} M'$ implies $\exists N'$ s.t $N \xRightarrow{\hat{\alpha}} N'$ and $M' \lesssim N'$

Theorem: Pre-congruence result

The binary relation \lesssim is a congruence over networks

Adapting tGNDC to wireless networks

Gorrieri and Martinelli's tGNDC is a general framework for the formal verification of security properties in a concurrent scenario. Intuitively:

A protocol M satisfies $tGNDC^{\rho(M)}$ if the presence on an arbitrary attacker does not affect M wrt the chosen abstraction $\rho(M)$ of the protocol.

tGNDC more formally:

A protocol M satisfies $tGNDC^{\rho(M)}$ if for any attacker A it holds that:

$$M \mid A \lesssim \rho(M)$$

Timed security properties:

By varying ρ we can express different timed security properties:

- **timed integrity**: freshness of authenticated packets
- **timed agreement**: agreement must be reached within a deadline

A sound proof technique for tGNDC

Proving that a protocol is tGNDC wrt some abstraction requires an universal quantification on all possible attackers. **The proof is hard!**

Definition: Top attacker

A^{TOP} denotes the Dolev-Yao attacker that can listen (and possibly replay) any message of the protocol. As usual it cannot guess secrets before they are disclosed

Theorem: Criterion for tGNDC

$$M \mid A^{\text{TOP}} \lesssim \rho(M) \text{ implies } M \mid A \lesssim \rho(M), \text{ for any } A$$

On the other hand, for proving that a protocol **is not tGNDC** it is sufficient to exhibit an attacker A and an **execution trace** for $M \mid A$ which cannot be mimicked by $\rho(M)$ (simulation semantics \subseteq trace semantics)

A case study: The LiSP protocol

- LiSP is a key management protocol for Wireless Sensor Networks
- A LiSP network consists of a *Key Server* (KS) and a set of *nodes* m_1, \dots, m_k
- The transmission time is split into time intervals Δ_{refresh} long
- The protocol employs two different key families:
 - *master keys* $k_{\text{KS}:m_j}$, one for each node m_j , for initial setup between m_j and BS
 - *temporal keys* k_0, \dots, k_n used by all nodes to encrypt/decrypt data packets
- Temporal key k_i is tied to time interval i and renewed every Δ_{refresh}
- At interval i , k_i is shared by all nodes and it is used for encryption

Our Security Analysis: key freshness

Timed integrity requirement for LiSP

- A node should authenticate only keys sent by KS in the last Δ_{refresh} time units
- In fact, if a node would authenticate an obsolete key (older than Δ_{refresh}) then it would not be synchronised with the rest of the network!

The LiSP specification (Key Server)

D_0	$\stackrel{\text{def}}{=} \sigma.D_1$	synchronise and move to D_1
D_i	$\stackrel{\text{def}}{=} [k_i \ k_{s+i} \vdash_{\text{enc}} t_i]$ $[\text{UpdateKey } t_i \vdash_{\text{pair}} u_i]$ $!\langle u_i \rangle.\sigma.\sigma.D_{i+1}$	for $i \geq 1$, encrypt k_{s+i} with k_i build the UpdateKey packet u_i broadcast r_i , and move to D_{i+1}
L_i	$\stackrel{\text{def}}{=} [?(r).l_{i+1}]\sigma.L_{i+1}$	wait for request packets
l_i	$\stackrel{\text{def}}{=} [r \vdash_{\text{fst}} r_1]l_i^1; \sigma.\sigma.L_i$	extract first component
l_i^1	$\stackrel{\text{def}}{=} [r_1 = \text{RequestKey}]l_i^2; \sigma.\sigma.L_i$	check if r_1 is a RequestKey
l_i^2	$\stackrel{\text{def}}{=} [r \vdash_{\text{snd}} m]$ $[k_{\text{KS}:m} \ k_{s+i} \vdash_{\text{enc}} w_i]$ $[k_{s+i} \vdash_{\text{hash}} h_i]$ $[w_i \ h_i \vdash_{\text{pair}} r_i]$ $[\text{InitKey } r_i \vdash_{\text{pair}} q_i]$ $\sigma.!\langle q_i \rangle.\sigma.L_i$	extract node name encrypt k_{s+i} with $k_{\text{KS}:m}$ calculate hash code for k_{s+i} build a pair r_i build a InitKey packet q_i broadcast q_i , move to L_i

The LiSP Protocol (receiver at node m)

Z	$\stackrel{\text{def}}{=}$	$[\text{RequestKey } m \vdash_{\text{pair}} r]$ $!\langle r \rangle.\sigma.[?(q).T]Z$	send a RequestKey packet wait for a reconfig. packet
T	$\stackrel{\text{def}}{=}$	$[q \vdash_{\text{fst}} q']T^1; \sigma.Z$	extract fst component of q
T^1	$\stackrel{\text{def}}{=}$	$[q' = \text{InitKey}]T^2; \sigma.Z$	check if q is a InitKey packet
T^2	$\stackrel{\text{def}}{=}$	$[q \vdash_{\text{snd}} q'']$ $[q'' \vdash_{\text{fst}} w]T^3; \sigma.Z$	extract snd component of q extract fst component of q''
T^3	$\stackrel{\text{def}}{=}$	$[q'' \vdash_{\text{snd}} h]$ $[k_{\text{KS}:m} w \vdash_{\text{dec}} k]T^3; \sigma.Z$	extract snd component of q'' extract the key
T^4	$\stackrel{\text{def}}{=}$	$[k \vdash_{\text{hash}} h'][h = h']T^5; \sigma.Z$	verify hash codes
T^5	$\stackrel{\text{def}}{=}$	$\sigma.\sigma.R\langle F^{s-1}(k), k, s-1 \rangle$	synchronise and move to R
$R(k_C, k_L, l)$	$\stackrel{\text{def}}{=}$	$[?(u).E]F$	wait for incoming packets
E	$\stackrel{\text{def}}{=}$	$[u \vdash_{\text{fst}} u']E^1; \sigma.F$	extract fst component of u
E^1	$\stackrel{\text{def}}{=}$	$[u' = \text{UpdateKey}]E^2; \sigma.F$	check UpdateKey packet
E^2	$\stackrel{\text{def}}{=}$	$[u \vdash_{\text{snd}} u'']$ $[k_C u'' \vdash_{\text{dec}} k]E^3; \sigma.F$	extract snd component of u decrypt u'' by using k_C
E^3	$\stackrel{\text{def}}{=}$	$[F^{s-l}(k) = k_L]E^4; \sigma.F$	authenticate k
E^4	$\stackrel{\text{def}}{=}$	$\sigma.\sigma.R\langle F^{s-1}(k), k, s-1 \rangle$	synchronise and move to R
F	$\stackrel{\text{def}}{=}$	$[l = 0]Z; \sigma.R\langle F^{l-1}(k_L), k_L, l-1 \rangle$	check if buffer key is empty

Verifying timed integrity 1/2

- The LiSP protocol, in its initial configuration, can be represented as:

$$\text{LiSP} \stackrel{\text{def}}{=} \prod_{j \in J} m_j[\sigma.Z]^{\nu_{m_j}} \mid \text{KL}[\sigma.L_0]^{\nu_{\text{KL}}} \mid \text{KD}[\sigma.D_0]^{\nu_{\text{KD}}}$$

where $m_j \in \nu_{\text{KD}} \cap \nu_{\text{KL}}$ and $\{\text{KD}, \text{KL}\} \subseteq \nu_{m_j}$

- For our analysis it is sufficient to consider only a part of it

$$\text{sLiSP} \stackrel{\text{def}}{=} m[\sigma.Z]^{\nu_m} \mid \text{KL}[\sigma.L_0]^{\nu_{\text{KL}}}$$

Definition: Timed integrity abstraction for sLiSP

$$\rho(\text{sLiSP}) \stackrel{\text{def}}{=} m[\sigma.\hat{Z}]^{\nu_m} \mid \text{KL}[\sigma.\hat{L}_0]^{\nu_{\text{KL}}}$$

for appropriate \hat{Z} and \hat{L}_0 .

Proposition: The abstraction is adequate

In $\rho(\text{sLiSP})$ key authentication occurs every Δ_{refresh} time units

Verifying timed integrity 2/2

Theorem: Replay attack to LiSP

There is an attacker A such that

$$\text{sLiSP} \mid A \not\sim \rho(\text{sLiSP}) .$$

Proof Give a trace of $\text{sLiSP} \mid A$ which cannot be matched by $\rho(\text{sLiSP})!$

$m \longrightarrow \text{KL} : r$	m sends a RequestKey to KL
$\text{KL} \longrightarrow m : q_1$	KL replies an InitKey lost by m and grasped by A
\longrightarrow	after Δ_{refresh} time units
$m \longrightarrow \text{KL} : r$	m sends a new RequestKey which gets lost
$A \longrightarrow m : q_1$	A replays the InitKey q_1 to m
\longrightarrow	after Δ_{refresh} time units
$m \rightarrow * : \text{auth}_1$	m authenticates the obsolete InitKey q_1

*** m has authenticated an InitKey which is $2\Delta_{\text{refresh}}$ old!!!***



Can we fix the problem?

Sure! By adding nonces in communications as in other security protocols

Let **nsLiSP** be the variant of sLiSP with nonces

Theorem: Timed integrity of nsLiSP

For any attacker A

$$\text{nsLiSP} \mid A \lesssim \rho(\text{nsLiSP}) .$$

Is the protocol with nonces safe now?

Well... when trying to prove *timed agreement* we found a different replay attack (for details see the full paper)

Conclusions

- We have proposed a process calculus to model wireless network security protocols
- The calculus comes with both an operational semantics and a simulation theory
- We have adapted Gorrieri and Martinelli's tGNDC to wireless systems
- Provided a soundness criterion for tGNDC
- Analysed the LiSP protocols and found a replay attack on key authentication
- and fixed the problem
- Can we use our technique to analyse other protocols? Yes, in the full paper we have applied our tGNDC to analyse both μ TESLA and *LEAP+* (here we found another replay attack)